arXiv:1603.09157v1 [stat.CO] 30 Mar 2016

Linear System Identification via EM with Latent Disturbances and Lagrangian Relaxation

Jack Umenberger* Johan Wågberg† Ian R. Manchester*
Thomas B. Schön†

August 7, 2016

• **Please cite this version:**

Jack Umenberger, Johan Wågberg, Ian R. Manchester and Thomas B. Schön. **Linear System Identification via EM with Latent Disturbances and Lagrangian Relaxation.** *arXiv, 2016.*

Abstract

In the application of the Expectation Maximization algorithm to identification of dynamical systems, internal *states* are typically chosen as latent variables, for simplicity. In this work, we propose a different choice of latent variables, namely, system *disturbances*. Such a formulation elegantly handles the problematic case of singular state space models, and is shown, under certain circumstances, to improve the fidelity of bounds on the likelihood, leading to convergence in fewer iterations. To access these benefits we develop a Lagrangian relaxation of the nonconvex optimization problems that arise in the latent disturbances formulation, and proceed via semidefinite programming.

Keywords: System identification, expectation maximization, Lagrangian relaxation, convex optimization.

*School of Aerospace, Mechanical and Mechatronic Engineering, University of Sydney, Australia e-mail: {j.umenberger, i.manchester}@acfr.usyd.edu.au

†Department of Information Technology, Uppsala University, Sweden, e-mail: {johan.wagberg, thomas.schon}@it.uu.se.

1 Introduction

Linear time invariant (LTI) state-space models provide a useful approximation of dynamical system behavior in a multitude of applications. In situations where models cannot be derived from first principles, some form of data-driven modeling, i.e. system identification, is appropriate [1]. This paper is concerned with identification of discrete-time LTI models of the form

$$x_{t+1} = Ax_t + Bu_t + Gw_t, \tag{1a}$$
$$y_t = Cx_t + Du_t + v_t, \tag{1b}$$

where $x_t \in \mathbb{R}^{n_x}$ denotes the system state, and $u_t \in \mathbb{R}^{n_u}$, $y_t \in \mathbb{R}^{n_y}$ denote the observed input and output, respectively. The disturbances (a.k.a. process noise), $w_t \in \mathbb{R}^{n_w}$ and measurement noise, v_t, are modeled as zero mean Gaussian white noise processes, while the uncertainty in the initial condition x_1 is modeled by a normal distribution, i.e.

$$w_t \sim \mathcal{N}(0, \Sigma_w),\ v_t \sim \mathcal{N}(0, \Sigma_v),\ x_1 \sim \mathcal{N}(\mu, \Sigma_1).$$

For convenience, all unknown model parameters are denoted by the variable $\theta = \{\mu, \Sigma_1, \Sigma_w, \Sigma_v, A, B, G, C, D\}$.

Despite the simplicity of LTI models, identification of such systems is complicated by the presence of *latent* variables. Specifically, in applications the observed data typically consists of inputs and (noisy) outputs, but not internal states or exogenous disturbances.

Various strategies have been developed to deal with this 'missing data'. *Marginalization*, for instance, involves integrating out (i.e. marginalizing over) the latent variables, leaving θ as the only quantity to be estimated. This is the approach adopted by prediction error methods [1, 2] and the Metropolis-Hastings algorithm [3, 4].

Alternatively, one may treat the latent variables as additional quantities to be estimated together with the model parameters. Such a strategy is termed *data augmentation*, and examples include subspace methods [5, 6], the Gibbs sampler [7], and the Expectation Maximization (EM) algorithm [8, 9, 10, 11].

Recently, a new family of methods have been developed in which one *supremizes over* the latent variables to obtain convex upper bounds for quality-of-fit cost functions, such as simulation error (a.k.a. output error) [12, 13, 14]. An important technique employed in this approach is Lagrangian relaxation [15, 16, 17], which replaces difficult constrained optimization problems with tractable, unconstrained, convex approximations.

This work draws on the underlying similarities between EM and Lagrangian relaxation to develop a new algorithm that seeks the maximum likelihood estimate of the model parameters θ, given measurements $u_{1:T}$ and $y_{1:T}$, i.e.

$$\theta^{\mathrm{ML}} = \arg\max_{\theta} p_\theta(u_{1:T}, y_{1:T}). \tag{2}$$

The EM algorithm is an iterative approach to ML estimation, in which estimates of the latent variables are used to construct tractable lower bounds to the likelihood. In the application of EM to (2) the latent variables are typically taken to be the system *states*, $x_{1:T}$, as this simplifies the ensuing optimization problem(s). Specifically, optimization of the bound at each iteration reduces to linear least squares [10].

In this work, we formulate the EM algorithm, for the approximate solution of (2), over latent *disturbances*, $w_{1:T}$. In contrast to the latent states formulation, the optimization of bounds based on disturbances is nonconvex. By applying Lagrangian relaxation, we obtain new bounds that can be optimized by semidefinite programming (SDP). The resulting algorithm can be considered an example of the more general *minorization maximization* principle [18, 19].

The rewards for this additional complexity are threefold. First, the proposed method elegantly handles identification of *singular* state-space models (i.e. $n_w < n_x$), a case to which the standard formulation of EM over latent states is not applicable, without modification. Secondly, this approach naturally ensures stability of the model at each iteration. Finally, when the *magnitude*

of the disturbances (i.e. Σ_w) is small, we show that use of latent disturbances produces better approximations to the likelihood, leading to convergence in fewer iterations.

We first introduced the basic idea of combining Lagrangian relaxation with a formulation of EM over latent disturbances in our conference paper [20]. This paper extends this recent work in several significant ways. In Section 4 we apply Lagrangian relaxation without resorting to Monte Carlo approximations, unlike the approach outlined in [20]. Furthermore, the Lagrangian relaxation detailed in this paper makes use of a more sophisticated multiplier, introduced in Section 4.4. A new study of the behavior of the EM algorithm for large and small disturbances is presented in Sections 5.2 and 6.1, offering insights into the results of numerical experiments on convergence rates in Section 6.2.

2 Preliminaries

2.1 Notation

The cone of real, symmetric nonnegative (positive) definite matrices is denoted by $\mathbb{S}^n_+$ ($\mathbb{S}^n_{++}$). The $n \times n$ identity matrix is denoted I_n. Let vec : $\mathbb{R}^{m\times n} \mapsto \mathbb{R}^{mn}$ denote the function that stacks the columns of a matrix to produce a column vector. The Kronecker product is denoted $\otimes$. The transpose of a matrix A is denoted A', and $|A|^2_Q$ is shorthand for $A'QA$. Time series data $\{x_t\}_{t=a}^b$ is denoted $x_{a:b}$ where $a, b \in \mathbb{N}$. A random variable x distributed according to the multivariate normal distribution, with mean μ and covariance Σ, is denoted $x \sim \mathcal{N}(x; \mu, \Sigma)$. We use $a(\theta) \propto b(\theta)$ to mean $b(\theta) = c_1 a(\theta) + c_2$ where c_1, c_2 are constants that do not effect optimization of $a(\theta)$ w.r.t. θ. For invertible A, $A\backslash B$ is shorthand for $A^{-1}B$. The log likelihood function is denoted $L_\theta(y_{1:T}) \triangleq \log p_\theta(u_{1:T}, y_{1:T})$.

2.2 The minorization-maximization principle

The minorization-maximization (MM) principle [18, 19] is an iterative approach to optimization problems of the form $\max_\theta\ f(\theta)$. Given an objective function $f(\theta)$ (not necessarily a likelihood), at each iteration of an MM algorithm we first build a *tight* lower bound $b(\theta, \theta_k)$ satisfying

$$f(\theta) \geq b(\theta, \theta_k)\ \forall\ \theta \text{ and } f(\theta_k) = b(\theta_k, \theta_k),$$

i.e. we *minorize* f by b. Then we optimize $b(\theta, \theta_k)$ w.r.t. θ to obtain θ_{k+1} such that $f(\theta_{k+1}) \geq f(\theta_k)$. The principle is useful when direct optimization of f is challenging, but optimization of b is tractable (e.g. concave). In the following two subsections, we present EM and Lagrangian relaxation as special cases of the MM principle, for problems involving missing data. Each of these algorithms is predicated on the assumption that there exists latent variables, z, such that optimization of $f(\theta)$ would be more straightforward if z were known.

2.3 The Expectation Maximization algorithm

The EM algorithm [8] applies the MM principle to ML estimation, i.e. $f(\theta) = L_\theta(y_{1:T})$. Each iteration of the algorithm consists of two steps: the expectation (E) step computes the *auxiliary* function

$$Q(\theta, \theta_k) \triangleq \int L_\theta(y_{1:T}, Z) p_{\theta_k}(Z \mid y_{1:T})\ \mathrm{d}Z = \mathrm{E}_{\theta_k}\left[L_\theta(y_{1:T}, Z) \,|\, y_{1:T}\right], \tag{3}$$

which is then maximized in lieu of the likelihood function during the maximization (M) step. The auxiliary function can be shown to satisfy the following inequality

$$L_\theta(y_{1:T}) - L_{\theta_k}(y_{1:T}) \geq Q(\theta, \theta_k) - Q(\theta_k, \theta_k) \tag{4}$$

and so the new parameter estimate θ_{k+1} obtained by maximization of $Q(\theta, \theta_k)$ is guaranteed to be of equal or greater likelihood than θ_k. In this sense, EM may be thought of as a specific MM

recipe for building lower bounds $Q(\theta, \theta_k)$ to the objective $L_\theta(y_{1:T})$, in ML estimation problems involving latent variables.

Remark 1. *Strictly speaking $Q(\theta, \theta_k)$ does not minorize $L_\theta(y_{1:T})$. Rather, the* change *in $Q(\theta, \theta_k)$ lower bounds the* change *in $L_\theta(y_{1:T})$; c.f. (4). Nevertheless, with some abuse of terminology, we will refer to $Q(\theta, \theta_k)$ as a lower bound, as shorthand for the relationship in (4).*

2.4 Lagrangian relaxation

The technique of Lagrangian relaxation applies the MM principle to *constrained* optimization problems of the form

$$\min_{\theta,z} \; J(\theta, z) \text{ s.t. } F(\theta, z) = 0, \tag{5}$$

i.e. $f(\theta) = J(\theta, z^*)$ where z^* is such that $F(\theta, z^*) = 0$. Here $J(\theta, z)$ is a cost function assumed to be convex in θ, and $F(\theta, z)$, assumed affine in θ, encodes the constraints. Notice that we present the problem as cost minimization, rather than objective maximization, and consequently develop upper bounds; however, this difference in superficial.

Unlike EM, in which we estimate z, Lagrangian relaxation *supremizes* over the latent variables to generate the bound. Specifically, the relaxation of (5) takes the form

$$\bar{J}_\lambda(\theta) = \sup_z \; J(\theta, z) + \lambda F(\theta, z), \tag{6}$$

where λ may be interpreted as a Lagrange multiplier. For arbitrary λ, the function $\bar{J}_\lambda(\theta)$ has two key properties:

1) It is convex in θ. Recall that J and F are convex and affine in θ, respectively. As such, $\bar{J}_\lambda(\theta)$ is the supremum of an infinite family of convex functions, and is, therefore, itself convex in θ; see Section 3.2.3 of [21].

2) It is an upper bound for the original problem (5). Given θ, let z^* be such that $F(\theta, z^*) = 0$. Then

$$J(\theta, z^*) + \lambda F(\theta, z^*) = J(\theta, z^*) = f(\theta),$$

which implies that the supremum over all z can be no smaller; i.e. $\bar{J}_\lambda(\theta)$ is an upper bound for $f(\theta)$.

The original optimization problem (5) may then be approximated by the convex program $\min_\theta \; \bar{J}_\lambda(\theta)$.

2.5 Latent variables for dynamical systems

In the application of EM to the identification of dynamical systems, there are two possible choices of latent variables: systems states, $x_{1:T}$, and initial conditions and disturbances $\{x_1, w_{1:T}\}$. Choosing latent states yields a joint likelihood function of the form

$$p_\theta(y_{1:T}, x_{1:T}) = \left[\prod_{t=1}^{T} p_\theta(y_t \mid x_t)\right]\left[\prod_{t=1}^{T-1} p_\theta(x_{t+1} \mid x_t)\right] p_\theta(x_1), \tag{7}$$

whereas latent disturbances leads to

$$p_\theta(y_{1:T}, x_1, w_{1:T}) = \left[\prod_{t=1}^{T} p_\theta(y_t \mid x_t)\right] p_\theta(w_{1:T}) p_\theta(x_1), \tag{8}$$

where $x_{t+1} = Ax_t + Bu_t + Gw_t$ for $t = 1, \ldots, T$. We denote this *simulated* state sequence by

$$\mathcal{X}_T(\theta, u_{1:T}, x_1, w_{1:T}) = \left\{x_{1:T} : x_t = A^{t-1}x_1 + \sum_{\tau=1}^{t-1} A^{t-1-\tau}(Bu_\tau + Gw_\tau)\right\}, \tag{9}$$

which, for given θ, is a *deterministic* mapping from initial conditions and disturbances to system states.

One can begin to understand the relationship between the choice of latent variables and difficulty of the ensuing optimization problems by examining the joint log likelihood. For latent states, $\log p_\theta(y_{1:T}, x_{1:T})$ decomposes as

$$\begin{aligned}\log p_\theta(y_{1:T}, x_{1:T}) = &\sum_{t=1}^{T} |y_t - Cx_t - Du_t|^2_{\Sigma_v^{-1}} + \sum_{t=1}^{T-1} |x_{t+1} - Ax_t - Bu_t|^2_{\Sigma_w^{-1}} \\ &+ |x_1 - \mu|^2_{\Sigma_1^{-1}} + T \log\det \Sigma_v + T \log\det \Sigma_w + \log\det \Sigma_1\end{aligned}$$

Given $\{u_{1:T}, x_{1:T}, y_{1:T}\}$, optimization of this function (w.r.t. θ) amounts to simple linear least squares. Conversely, for latent disturbances, $\log p_\theta(y_{1:T}, x_1, w_{1:T})$, is given by

$$\begin{aligned}\log p_\theta(y_{1:T}, x_1, w_{1:T}) = &\mathcal{E}(\theta, u_{1:T}, y_{1:T}, x_1, w_{1:T}) + \sum_{t=1}^{T} |w_t|^2_{\Sigma_w^{-1}} + |x_1 - \mu|^2_{\Sigma_1^{-1}} \\ &+ T \log\det \Sigma_v + T \log\det \Sigma_w + \log\det \Sigma_1.\end{aligned} \tag{10}$$

Here $\mathcal{E}(\theta, u_{1:T}, y_{1:T}, x_1, w_{1:T})$ denotes the *simulation error*, defined

$$\mathcal{E}(\theta, u_{1:T}, y_{1:T}, x_1, w_{1:T}) \triangleq \sum_{t=1}^{T} |y_t - Cx_t - Du_t|^2_{\Sigma_v^{-1}} \tag{11}$$

where $x_{1:T} = \mathcal{X}_T(\theta, u_{1:T}, x_1, w_{1:T})$. This dependence on the simulated state sequence renders optimization of (10) a challenging nonlinear, nonconvex problem.

3 EM with latent disturbances

In this section we detail the application of EM to the identification of LGSS models, when formulated with latent disturbances; refer to [10] for the formulation over latent states. Each iteration of the algorithm involves optimization of the *auxiliary function*

$$\begin{aligned}Q(\theta, \theta_k) &= \mathrm{E}_{\theta_k}\left[\log p_\theta(y_{1:T}, x_1, w_{1:T}) \,|\, y_{1:T}\right] \\ &= \int \log p_\theta(y_{1:T}, x_1, w_{1:T}) p_{\theta_k}(x_1, w_{1:T} \mid y_{1:T}) \,\mathrm{d}x_1 \mathrm{d}w_{1:T}\end{aligned} \tag{12}$$

which serves as a lower bound to the likelihood, given our current best estimate of the model parameters, θ_k. With the joint log likelihood $\log p_\theta(y_{1:T}, x_1, w_{1:T})$ given by (8), to evaluate $Q(\theta, \theta_k)$ we must compute $p_{\theta_k}(x_1, w_{1:T} \mid y_{1:T})$, i.e., the joint smoothing distribution (JSD) of the initial state and disturbances. The E step then amounts to solving a *disturbance smoothing*, rather than *state smoothing*, problem, which reflects the use of disturbances, rather than states, as latent variables.

As demonstrated in Section 2.5, $L_\theta(y_{1:T}, x_1, w_{1:T})$ involves the simulated state sequence $\mathcal{X}_T(\theta, u_{1:T}, x_1, w_{1:T})$. As a consequence, we shall show that the M step is equivalent to a nonconvex simulation error minimization problem (c.f. (23)). This is in contrast to the latent states formulation, in which $\max_\theta \; Q(\theta, \theta_k)$ reduces to linear least squares.

3.1 Expectation step

To compute $Q(\theta, \theta_k)$ it is convenient to use the following decomposition

$$\begin{aligned}Q(\theta, \theta_k) = &\underbrace{\mathrm{E}_{\theta_k}\left[\log p_\theta(x_1) \,|\, y_{1:T}\right]}_{Q_1(\theta, \theta_k)} + \underbrace{\mathrm{E}_{\theta_k}\left[\log p_\theta(w_{1:T}) \,|\, y_{1:T}\right]}_{Q_2(\theta, \theta_k)} \\ &+ \underbrace{\mathrm{E}_{\theta_k}\left[\log p_\theta(y_{1:T} \mid x_1, w_{1:T}) \,|\, y_{1:T}\right]}_{Q_3(\theta, \theta_k)}\end{aligned} \tag{13}$$

which was obtained by inserting (8) into (12).

Remark 2. *Each term in (13) is a function of different parameters: μ and Σ_1 appear only in $Q_1(\theta, \theta_k)$; Σ_w in $Q_2(\theta, \theta_k)$; and Σ_v, A, B, G, C, D in $Q_3(\theta, \theta_k)$. To emphasize this we introduce the following decomposition of θ*

$$\alpha = \{\mu, \Sigma_1\},\ \beta = \Sigma_w,\ \gamma = \{\Sigma_v, A, B, G, C, D\}.$$

The following lemma details the computation of $Q(\theta, \theta_k)$. For clarity of expression, we introduce the following *lifted* form of the dynamics in (1),

$$Y = \bar{C}\bar{F}Z + (\bar{C}\bar{G} + \bar{D})U + V$$

where $Y = \mathrm{vec}(y_{1:T})$, $U = \mathrm{vec}(u_{1:T})$, $V = \mathrm{vec}(v_{1:T})$, $Z = \mathrm{vec}([x_1, w_{1:T-1}])$,

$$\bar{F} = \begin{bmatrix} I & 0 & 0 & 0 & \dots & 0 \\ A & G & 0 & 0 & \dots & 0 \\ A^2 & A & G & 0 & \dots & 0 \\ \vdots & & & & \ddots & \vdots \\ A^{T-1} & A^{T-2}G & A^{T-3}G & & \dots & G \end{bmatrix},$$

$$\bar{G} = \begin{bmatrix} 0 & 0 & 0 & 0 & \dots & 0 \\ B & 0 & 0 & 0 & \dots & 0 \\ AB & B & 0 & 0 & \dots & 0 \\ \vdots & & & & \ddots & \vdots \\ A^{T-2}B & A^{T-3}B & \dots & AB & B & 0 \end{bmatrix},$$

$\bar{C} = I_T \otimes C$ and $\bar{D} = I_T \otimes D$.

Lemma 1. *The auxiliary function $Q(\theta, \theta_k)$ defined in (12) is given by*

$$\begin{aligned} Q(\theta, \theta_k) \propto &- \log\det\Sigma_1 - |\hat{x}_{1|T} - \mu|^2_{\Sigma_1^{-1}} - \mathrm{tr}\big(\Sigma_1^{-1}\widehat{\Sigma}_{1|T}\big) \\ &- T\log\det\Sigma_w - \sum_{t=1}^{T} \mathrm{tr}\big(\Sigma_w^{-1}\,\mathrm{E}_{\theta_k}\left[w_t w_t' \,|\, Y\right]\big) \\ &- T\log\det\Sigma_v - \mathrm{tr}(\Sigma_Y^{-1}(\bar{C}\bar{F}\Omega\bar{F}'\bar{C}' + \Delta\Delta')) \end{aligned}$$

where

$$\hat{x}_{1|T} = \mathrm{E}_{\theta_k}\left[x_1 \,|\, y_{1:T}\right], \tag{14a}$$
$$\widehat{\Sigma}_{1|T} = \mathrm{Var}_{\theta_k}\left[x_1 \,|\, y_{1:T}\right], \tag{14b}$$

$$\hat{Z} = \mathrm{E}_{\theta_k}\left[Z \mid y_{1:T}\right], \tag{15a}$$
$$\Omega = \mathrm{Var}_{\theta_k}\left[Z \mid y_{1:T}\right], \tag{15b}$$

$$\mu_Y \triangleq \mathrm{E}_{\theta}\left[Y \,|\, Z\right] = \bar{C}\bar{F}Z + (\bar{C}\bar{G} + \bar{D})U, \tag{16a}$$
$$\Sigma_Y \triangleq \mathrm{Var}_{\theta}\left[Y \,|\, Z\right] = I_T \otimes \Sigma_v, \tag{16b}$$

$$\Delta = \mathrm{E}_{\theta_k}\left[Y - \mu_Y \,|\, y_{1:T}\right] = Y - \bar{C}\bar{F}\hat{Z} - (\bar{C}\bar{G} + \bar{D})U. \tag{17}$$

Proof. The first term in (13) is given by

$$\begin{aligned} \mathrm{E}_{\theta_k}\left[\log p_\theta(x_1) \,|\, y_{1:T}\right] &= \mathrm{E}_{\theta_k}\left[\log\mathcal{N}(x_1; \mu, \Sigma_1) \,|\, y_{1:T}\right] \\ &= \mathrm{E}_{\theta_k}\left[-\frac{n_x}{2}\log 2\pi - \frac{1}{2}\log\det\Sigma_1 - |x_1 - \mu|^2_{\Sigma_1^{-1}} \,\middle|\, y_{1:T}\right]. \end{aligned}$$

Ignoring constant terms and scaling factors yields

$$Q_1(\alpha, \theta_k) \propto -\log\det\Sigma_1 - |\hat{x}_{1|T} - \mu|^2_{\Sigma_1^{-1}} - \mathrm{tr}\big(\Sigma_1^{-1}\widehat{\Sigma}_{1|T}\big) \tag{18}$$

where $\hat{x}_{1|T}$ and $\widehat{\Sigma}_{1|T}$ are given in (14).

As the disturbances are i.i.d., $Q_2(\beta, \theta_k)$ is given by

$$\begin{aligned}\mathrm{E}_{\theta_k}\left[\log p_\theta(w_{1:T}) \,|\, y_{1:T}\right] &= \mathrm{E}_{\theta_k}\left[\log \prod_{t=1}^{T} \mathcal{N}(w_t; 0, \Sigma_w) \,\middle|\, y_{1:T}\right] \\ &= \sum_{t=1}^{T} \mathrm{E}_{\theta_k}\left[-\frac{n_w}{2}\log 2\pi - \frac{1}{2}\log\det\Sigma_w - \frac{1}{2}|w_t|^2_{\Sigma_w^{-1}} \,\middle|\, y_{1:T}\right].\end{aligned}$$

Once more ignoring constants, this reduces to

$$Q_2(\beta, \theta_k) \propto -T\log\det\Sigma_w - \sum_{t=1}^{T} \mathrm{tr}\big(\Sigma_w^{-1}\,\mathrm{E}_{\theta_k}\left[w_t w_t' \,|\, y_{1:T}\right]\big). \tag{19}$$

Finally, we turn our attention to $Q_3(\gamma, \theta_k)$. The p.d.f. $p_\theta(y_{1:T} \mid x_1, w_{1:T})$ is given by $p_\theta(Y \mid Z) = \mathcal{N}(Y; \mu_Y, \Sigma_Y)$, where μ_Y and Σ_Y are given in (16). $Q_3(\gamma, \theta_k)$ may then be expressed as

$$\mathrm{E}_{\theta_k}\left[\log\mathcal{N}(Y; \mu_Y, \Sigma_Y) \,|\, y_{1:T}\right] = -\frac{Tn_y}{2}\log 2\pi - \log\det\Sigma_Y - \mathrm{E}_{\theta_k}\left[|Y-\mu_Y|^2_{\Sigma_Y^{-1}} \,\middle|\, y_{1:T}\right].$$

Letting $\hat{Z}$ and Ω, defined in (15), denote the mean and covariance (respectively) of $p_{\theta_k}(x_1, w_{1:T-1}|y_{1:T})$, gives

$$Q_3(\gamma, \theta_k) \propto -T\log\det\Sigma_v - \mathrm{tr}(\Sigma_Y^{-1}(\bar{C}\bar{F}\Omega\bar{F}'\bar{C}' + \Delta\Delta')) \tag{20}$$

where $\Delta = \mathrm{E}_{\theta_k}\left[Y - \mu_Y \,|\, y_{1:T}\right]$ is defined in (17). ■

Calculating the quantities in (14) amounts to a state smoothing problem, the solution for which is given in closed form by, e.g., the RTS smoother [22] (see also, [23, Section 4.4]). Similarly, for the LGSS models considered in this work, $\mathrm{E}_{\theta_k}\left[w_t w_t' \,|\, Y_T\right]$, $\hat{Z}$ and Ω can be computed in closed form by standard disturbance smoothers; see, e.g., [23, Section 4.5].

3.2 Maximization step

To perform the M step, i.e. maximize $Q(\theta, \theta_k)$, we will utilize the same decomposition as in (13), and optimize each of the conditional expectations separately; the validity of this approach is established by Remark 2.

We begin with maximization of $Q_1(\alpha, \theta_k)$:
Lemma 2. *The solution to* $\alpha_{k+1} = \arg\max_\alpha\ Q_1(\alpha, \theta_k)$ *is given by* $\alpha_{k+1} = \{\hat{x}_{1|T}, \widehat{\Sigma}_{1|T}\}$.

Proof. To maximize $Q_1(\alpha, \theta_k)$, notice that (18) is concave w.r.t. μ and Σ_1^{-1}. Therefore, setting the gradient to zero gives the global maximizers $\mu = \hat{x}_{1|T}$ and $\Sigma_1 = \widehat{\Sigma}_{1|T}$. ■

Maximization of $Q_2(\beta, \theta_k)$ can be handled in a similar way:
Lemma 3. *The solution to* $\beta_{k+1} = \arg\max_\beta\ Q_2(\beta, \theta_k)$ *is given by* $\beta_{k+1} = \hat{\Sigma}_w$ *where*

$$\hat{\Sigma}_w = \frac{1}{T}\sum_{t=1}^{T} \mathrm{E}_{\theta_k}\left[w_t w_t' \,|\, y_{1:T}\right]. \tag{21}$$

Proof. Substituting (21) into (19) yields

$$Q_2(\beta, \theta_k) \propto -T\log\det\Sigma_w - T\mathrm{tr}\big(\Sigma_w^{-1}\hat{\Sigma}_w\big).$$

This function is concave w.r.t. Σ_w and so setting the gradient to zero gives the global maximizer $\Sigma_w = \hat{\Sigma}_w$. ■

Finally, we consider maximization of $Q_3(\gamma, \theta_k)$. This is a challenging problem, due to its dependence on *simulated* state sequences; c.f. Section 2.5. Indeed, from (20), it is clear that the quantities $\bar{F}$ and $\bar{G}$ render $Q_3(\gamma, \theta_k)$ a nonconvex function of the model parameters.

To maximize $Q_3(\gamma, \theta_k)$ it is convenient to conceptualize (20) as the summation of $T+1$ simultaneous simulation error minimization problems.

Lemma 4. *Recalling the definition of simulation error in (11), $Q_3(\gamma, \theta_k)$ in (13) is equivalent to:*

$$-Q_3(\gamma, \theta_k) = \mathcal{E}(\theta, u_{1:T}, y_{1:T}, \hat{x}_{1|T}, \hat{w}_{1:T}) + \sum_{j=1}^{T} \mathcal{E}(\theta, 0, 0, x_1^j, w_{1:T}^j) + T \log\det \Sigma_v \tag{22}$$

where $x_1^j, w_{1:T}^j$ are such that $\Omega = \sum_{j=1}^{T} \omega_j \omega_j'$ for $\omega_j = \mathrm{vec}([x_1^j, w_{1:T-1}^j])$.

Proof. First consider the $\mathrm{tr}(\Sigma_Y \Delta\Delta')$ term in (20). From (17), Δ is clearly the difference between the measured output $y_{1:T}$ and the simulated output of the model with the expected value of the latent disturbances, i.e.

$$\{\hat{x}_{1|T}, \hat{w}_{1:T-1}\} = \mathrm{E}_{\theta_k}\left[x_1, w_{1:T-1} \mid y_{1:T}\right] = \hat{Z}.$$

Therefore,

$$\mathrm{tr}(\Sigma_Y \Delta\Delta') = \sum_{t=1}^{T} |y_t - Cx_t - Du_t|^2_{\Sigma_v^{-1}}$$

where $x_{t+1} = Ax_t + Bu_t + G\hat{w}_t$ with $x_1 = \hat{x}_{1|T}$.

Next, consider the $\mathrm{tr}(\Sigma_Y^{-1} \bar{C}\bar{F}\Omega\bar{F}'\bar{C}')$ term. Decomposing Ω as the sum of T rank one matrices, i.e. $\Omega = \sum_{j=1}^{T} \omega_j \omega_j'$, leads to

$$\mathrm{tr}(\Sigma_Y^{-1} \bar{C}\bar{F}\Omega\bar{F}'\bar{C}') = \sum_{j=1}^{T} |\bar{C}\bar{F}w_j|^2_{\Sigma_Y^{-1}} = \sum_{j=1}^{T}\sum_{t=1}^{T} |Cx_t^j|^2_{\Sigma_v^{-1}}$$

where $x_{t+1}^j = Ax_t^j + Gw_t^j$. One can interpret this as the sum of T simulation error problems with $y_{1:T} \equiv 0$, $u_{1:T} \equiv 0$, $\{x_1, w_{1:T}\} = \omega_j$. ∎

To summarize, the computations involved in each iteration of the EM algorithm (formulated with latent disturbances) are straightforward, with the exception of maximization of $Q_3(\gamma, \theta_k)$. From (22), this maximization is equivalent to $T+1$ simultaneous nonconvex simulation error minimization problems.

4 Lagrangian relaxation of maximization step

In this section we describe how Lagrangian relaxation (c.f. Section 2.4) can be applied to the optimization of $Q_3(\gamma, \theta_k)$ in (22). Specifically, we shall develop a bound for $Q_3(\gamma, \theta_k)$ that can be efficiently optimized as a convex program. Furthermore, we shall show how this approach naturally enforces model stability at each iteration, by searching over a convex parametrization of all stable linear models.

4.1 Lagrangian relaxation of simulation error

Optimization of $Q_3(\gamma, \theta_k)$ in (22) is difficult because it requires minimization of *simulation error*, as defined in (11). In this subsection, we detail the application of Lagrangian relaxation, introduced in Section 2.4, to minimization of simulation error, which can be formulated as

$$J^* \triangleq \min_{\gamma, x_{1:T}} J(\gamma, x_{1:T}) \triangleq \sum_{t=1}^{T} |y_t - Cx_t - Du_t|^2_{\Sigma_v^{-1}} \tag{23a}$$

$$\text{s.t. } \mathcal{F}(\gamma, u_{1:T}, \xi_1, x_{1:T}, w_{1:T}) = 0. \tag{23b}$$

Here ξ_1 denotes the initial state, assumed to be known, and $\mathcal{F}(\gamma, u_{1:T}, \xi_1, x_{1:T}, w_{1:T})$ encodes the dynamic constraints on $x_{1:T}$, such that

$$\mathcal{F}(\gamma, u_{1:T}, \xi_1, \mathcal{X}_T(\theta, u_{1:T}, \xi_1, w_{1:T}), w_{1:T}) = 0.$$

As in (6), the Lagrangian relaxation of (23) takes the form

$$\bar{J}_\lambda(\gamma, u_{1:T}, y_{1:T}, \xi_1, w_{1:T}) \triangleq \sup_{x_{1:T}} \{J(\gamma, x_{1:T}) - \lambda' \mathcal{F}(\gamma, u_{1:T}, \xi_1, x_{1:T}, w_{1:T})\}, \tag{24}$$

and, for arbitrary multiplier λ, represents a convex upper bound on the simulation error $\mathcal{E}(\gamma, u_{1:T}, y_{1:T}, x_1, w_{1:T})$.

4.2 Implicit dynamics

It remains to choose the Lagrange multiplier λ such that $\bar{J}_\lambda(\gamma)$ is a useful upper bound, i.e., such that $\bar{J}^*_\lambda \approx J^*$. Unfortunately, the simultaneous search for λ and γ is not jointly convex, due to the coupling between λ and $\mathcal{F}$, and so λ must be specified in advance. However, we can alleviate this restriction by searching over an implicit representation of the dynamics in (1a)

$$Ex_{t+1} = Fx_t + Ku_t + Lw_t, \tag{25a}$$
$$y_t = Cx_t + Du_t + v_t, \tag{25b}$$

where E is invertible such that $A = E^{-1}F$, $B = E^{-1}K$ and $G = E^{-1}L$. With the implicit dynamics of (25), the dynamics constraint can be expressed

$$\mathcal{F}(\eta, u_{1:T}, \xi_1, x_{1:T}, w_{1:T}) = \bar{F}\mathrm{vec}(x_{1:T}) + \epsilon \tag{26}$$

where $\bar{F} \in \mathbb{R}^{Tn_x \times Tn_x}$ and $\epsilon \in \mathbb{R}^{Tn_x}$ denote

$$\begin{bmatrix} E & 0 & \dots & & \\ -F & E & 0 & & \\ 0 & -F & E & 0 & \\ \vdots & & & \ddots & \ddots \end{bmatrix} \;\&\; \begin{bmatrix} -E\xi_1 \\ Ku_1 + Lw_1 \\ \vdots \\ Ku_{T-1} + Lw_{T-1} \end{bmatrix}$$

respectively. One may interpret the convex bound resulting from this implicit formulation as that of (24), but with the multiplier $(I \otimes E')\lambda$, thereby allowing a simultaneous (partial) search for multipliers and model parameters.

Remark 3. *We introduce $\eta = \{E, F, K, L, C, D, \Sigma_v, P\}$ to group the implicit model parameters, Σ_v and $P \in \mathbb{S}^{n_x}_{++}$, into a single variable. Here P represents a model stability certificate, the role of which is made precise in Lemma 6. Henceforth, $\bar{J}_\lambda(\eta)$ denotes Lagrangian relaxation with the implicit dynamics constraint (26). For convenience, we define the mapping $\mathcal{M} : \eta \mapsto \gamma$ from an implicit to explicit parametrization: $\mathcal{M}(\eta) \triangleq \{\Sigma_v, E\backslash F, E\backslash K, E\backslash L, C, D\}$.*

4.3 Convex upper bound for $-Q_3(\gamma, \theta_k)$

The representation of $Q_3(\gamma, \theta_k)$ in (22) makes the application of Lagrangian relaxation straightforward. To obtain a convex upper bound for $-Q_3(\gamma, \theta_k)$ we can simply replace each simulation error term $\mathcal{E}(\gamma)$ with the appropriate corresponding convex bound $\bar{J}_\lambda(\eta)$.

Lemma 5. *Consider the following function*

$$\begin{aligned} \bar{Q}_3(\eta) \triangleq\; & \bar{J}_{\lambda^0}(\eta, u_{1:T}, y_{1:T}, \hat{x}_{1|T}, \hat{w}_{1:T}) + \sum_{j=1}^{T} \bar{J}_{\lambda^j}(\eta, 0, 0, x_1^j, w_{1:T}^j) \\ & + T\mathrm{tr}\Sigma_{v_k}^{-1}\Sigma_v + T\log\det\Sigma_{v_k} + Tn_y, \end{aligned} \tag{27}$$

where $x_1^j, w_{1:T}^j$ are defined in Lemma 4. $\bar{Q}_3(\eta)$ is a convex upper bound for $-Q_3(\gamma, \theta_k)$, where $\gamma = \mathcal{M}(\eta)$.

Proof. As $\bar{Q}_3(\eta)$ is defined by a summation of convex functions, it is itself a convex function. Summation of the following inequalities

$$\begin{aligned}
\bar{J}_{\lambda^0}(\eta, u_{1:T}, y_{1:T}, \hat{x}_{1|T}, \hat{w}_{1:T}) &\geq \mathcal{E}(\gamma, u_{1:T}, y_{1:T}, \hat{x}_{1|T}, \hat{w}_{1:T}), \\
\bar{J}_{\lambda^j}(\eta, 0, 0, x_1^j, w_{1:T}^j) &\geq \mathcal{E}(\gamma, 0, 0, x_1^j, w_{1:T}^j), \quad j = 1, \ldots, T, \\
\operatorname{tr}(\Sigma_{v_k}^{-1}\Sigma_v) + \log\det\Sigma_{v_k} + n_y &\geq \log\det\Sigma_v,
\end{aligned}$$

gives $\bar{Q}_3(\eta) \geq -Q_3(\gamma, \theta_k)$. Notice that $n_y + \log\det\Sigma_{v_k} + \operatorname{tr}(\Sigma_{v_k}^{-1}\Sigma_v)$ is an affine upper bound on the concave term $\log\det\Sigma_v$, which is tight at our current best estimate of the covariance, Σ_{v_k}. ■

Notice that $\bar{Q}_3(\eta)$ is a function of an implicit representation of the dynamical model, denoted η, reflecting the fact that we formulate the Lagrangian relaxation using the implicit dynamics of (25a).

4.4 Lagrange multipliers

To utilize the convex bound $\bar{Q}_3(\eta)$ we must supply suitable Lagrange multipliers, $\{\lambda^j\}_{j=0}^T$. While convexity of the upper bound $\bar{J}_\lambda$ defined in (24) is guaranteed for any multiplier λ that is independent of η, in this work we consider multipliers of the form $\lambda = \operatorname{vec}\left(\{\lambda_t\}_{t=1}^T\right)$ for $\lambda_t = 2\left(Hx_t + h_t\right)$, i.e.

$$\lambda = 2\left(\Lambda\operatorname{vec}(x_{1:T}) + h\right), \tag{28}$$

where $\Lambda = I_T \otimes H$ for $H \in \mathbb{R}^{n_x \times n_x}$ and $h \in \mathbb{R}^{Tn_x}$. Recall from Section 4.2 that the use of the implicit model class (25) allows a convex (partial) search over model parameters and multipliers. Furthermore, this implicit representation permits the following definition of a convex parametrization of all stable LTI models.

Lemma 6. *Let $\Theta(H)$ denote the set of all models η of the form (25) and $P \in \mathbb{S}_{++}^{n_x}$ that satisfy the LMI*

$$M(\eta, H) = \begin{bmatrix} H'E + E'H - P & F'H & C' \\ H'F & P & 0 \\ C & 0 & \Sigma_v \end{bmatrix} > 0 \tag{29}$$

i.e. $\Theta(H) \triangleq \{\eta : M(\eta, H) > 0\}$.

Then a model θ of the form (1) is stable iff there exists E such that $\eta = \{E, EA, EB, EG, C, D, P\} \in \Theta(H)$ for some full-rank $H \in \mathbb{R}^{n_x \times n_x}$.

Proof. This result is a straightforward extension of Lemma 4 and Corollary 5 in [24, Section 3.2]. ■

Remark 4. *The LMI $M(\eta, H) > 0$ implies $H'E + E'H > 0$ which ensures that E is invertible, i.e., the implicit dynamics in (25) are well-posed.*

The model stability condition (29) and multiplier (28) also guarantee finiteness of the supremum in (24):

Lemma 7. *Given arbitrary $u_{1:T}$, $y_{1:T}$, ξ_1, $w_{1:T}$, h and full-rank H the supremum in the definition of $\bar{J}_\lambda(\eta)$ given by (24) is finite, for λ given by (28) and $\eta \in \Theta(H)$.*

Proof. For ease of exposition, we define

$$J_\lambda(\eta, x_{1:T}) \triangleq J(\eta, x_{1:T}) - \lambda'\mathcal{F}(\eta, u_{1:T}, \xi_1, x_{1:T}, w_{1:T}), \tag{30}$$

where $u_{1:T}$, $w_{1:T}$, ξ_1 are dropped from the notation for brevity. The bound in (24) may then be equivalently expressed as $\bar{J}_\lambda(\eta) = \sup_{x_{1:T}} J_\lambda(\eta, x_{1:T})$. We can write

$$J_\lambda(\eta, x_{1:T}) = \sum_{t=1}^{T} x_t' C'\Sigma_v^{-1} C x_t - 2\sum_{t=1}^{T-1}(Hx_{t+1})'(Ex_{t+1} - Fx_t) - 2(Hx_1)'Ex_1 + \operatorname{aff}(x_{1:T}) \tag{31}$$

where $\text{aff}(x_{1:T})$ denotes additional terms that are affine in $x_{1:T}$. We make use of the inequality

$$2a'b \leq |a|_P^2 + |b|_{P^{-1}}^2 \quad \forall\ a, b, P > 0$$

(see, e.g., [12, Section IV]) to obtain an upper bound for $J_\lambda(\eta, x_{1:T})$. Specifically, by setting $a = x_{t+1}$ and $b = H'Fx_t$ we obtain the inequality

$$2x_{t+1}'H'Fx_t \leq |x_{t+1}|_P^2 + |H'Fx_t|_{P^{-1}}^2 \tag{32}$$

which holds for all x_t, x_{t+1}. Applying to (32) to (31) yields the following upper bound:

$$\begin{aligned} J_\lambda(\eta, x_{1:T}) \leq\ & x_1' \left(|H'F|_{P^{-1}}^2 + C'\Sigma_v^{-1}C\right) x_1 \\ & + \sum_{t=2}^{T-1} x_t' \left(|H'F|_{P^{-1}}^2 + P - 2H'E + |C|_{\Sigma_v^{-1}}^2\right) x_t \\ & + x_T' \left(P - 2H'E + |C|_{\Sigma_v^{-1}}^2\right) x_T + \text{aff}(x_1 : T). \end{aligned} \tag{33}$$

The supremum w.r.t. $x_{1:T}$ of the upper bound on the LHS of (33) is finite when the quadratic component is concave, i.e., $|H'F|_{P^{-1}}^2 + P - H'E - E'H + |C|_{\Sigma_v^{-1}}^2 < 0$. By the Schur complement this condition is equivalent to the LMI in (29). As finiteness of the bound implies finiteness of $J_\lambda(\eta, x_{1:T})$ this completes the proof. ■

The key to the EM algorithm is (4), i.e., increasing $Q(\theta, \theta_k)$ guarantees an improvement in $L_\theta(y_{1:T})$. Consequently, we must ensure that optimization of $\bar{Q}_3(\eta)$ does not decrease $Q_3(\gamma, \theta_k)$. This property holds if there exist multipliers $\{\lambda^j\}_{j=0}^T$ such that the bound $\bar{Q}_3(\eta)$ is 'tight' to $-Q_3(\gamma, \theta_k)$ at $\gamma = \mathcal{M}(\eta) = \gamma_k$, and may be understood as an application of the MM principle of Section 2.3.

To obtain such a set of multipliers, we can minimize the bound $\bar{Q}_3(\eta_k)$ w.r.t. the multipliers $\{\lambda^j\}_{j=1}^T$ for a fixed η_k. Here η_k is such that $\gamma_k = \mathcal{M}(\eta_k)$. We propose a two-stage approach:

i. For each of the $T+1$ bounds $\bar{J}_{\lambda^j}(\eta)$ that comprise $\bar{Q}_3(\eta)$, solve the convex optimization problem

$$H_j = \arg\min_{\Phi_j}\ \bar{J}_{\phi_j}(\eta_k) \text{ s.t. } \eta_k \in \Theta(\Phi_j), \tag{34}$$

where $\phi_j = I_T \otimes \Phi_j \text{vec}(x_{1:T})$.

ii. Set $\lambda^j = 2\left(I_T \otimes H_j \text{vec}(x_{1:T}) + h_j\right)$ such that $\bar{J}_{\lambda^j}(\eta_k) = \mathcal{E}(\gamma_k)$, where h_j is computed as in Lemma 8.

Lemma 8. *Given a model $\eta_k \in \Theta(H)$ of the form (25), and arbitrary $u_{1:T}$, $y_{1:T}$, x_1, $w_{1:T}$ and full-rank H, let λ denote a multiplier of the form (28) with h defined as*

$$h = (\bar{F}')^{-1}\left(\Psi X^* - \bar{C}'\bar{\Sigma}_Y^{-1}(Y - \bar{D}U) + \Lambda'\epsilon\right). \tag{35}$$

Furthermore, let θ_k be such that $A = E\backslash F$, $B = E\backslash K$ and $G = E\backslash L$. With this multiplier $\bar{J}_\lambda(\eta_k) = \mathcal{E}(\theta_k)$, i.e. the convex bound $\bar{J}_\lambda(\eta)$ is tight to the simulation error $\mathcal{E}(\gamma)$ at $\eta = \eta_k$. The notation is as follows: $\bar{F}$ and ϵ are defined in (26); $\bar{C}$, $\bar{D}$, Σ_Y, U, Y are defined in Section 3.1; $X^* = \text{vec}(\mathcal{X}_T(\eta_k, u_{1:T}, x_1, w_{1:T}))$ and

$$\Psi = \bar{C}'\bar{\Sigma}_Y^{-1}\bar{C} - \Lambda'\bar{F} - \bar{F}'\Lambda.$$

Proof. As $\eta_k \in \Theta$, by Lemma 7, we know that $J_\lambda(\eta, x_{1:T})$, defined in (30), is a concave quadratic function in $x_{1:T}$. By the first order optimality condition, it can be shown that the state sequence $x_{1:T}^*$ that maximizes this function must satisfy

$$\Psi \text{vec}(x_{1:T}^*) = \bar{C}'(Y - \bar{D}U) + \bar{F}'h - \Lambda'\epsilon.$$

Setting $x_{1:T}^* = X^*$ and solving for h yields the expression in (35). Note that invertibility of E ensures that $\bar{F}^{-1}$ is well-defined; c.f. Remark 4. As $x_{1:T}^* = X^*$ we have

$$\bar{J}_\lambda(\eta_k) = J(\eta_k, X^*) - \lambda\mathcal{F}(\eta_k, u_{1:T}, X^*, w_{1:T}) = J(\eta_k, X^*) = \mathcal{E}(\eta_k).$$

■

In summary, to update γ at the k^{th} iteration of the EM algorithm, we solve the convex optimization problem

$$\eta_{k+1} = \arg\min_{\eta} \ \bar{Q}_3(\eta) \text{ s.t. } \eta \in \bigcap_{j=0}^{T} \Theta(H_j), \tag{36}$$

where λ^j is given by (28) with H_j from (34) and h_j from (35), for $j = 0, \ldots, T$, and then set $\gamma_{k+1} = \mathcal{M}(\eta_{k+1})$. For a complete summary of EM with latent disturbances, refer to Algorithm 1.

Remark 5. *A common heuristic for terminating the EM algorithm is to cease iterations once the change in likelihood falls below a certain tolerance δ, i.e.*

$$L_{\theta_{k+1}}(y_{1:T}) - L_{\theta_k}(y_{1:T}) < \delta. \tag{37}$$

Alternatively, one can simply run the algorithm for a finite number of iterations, chosen so as to attain a model of quality sufficient for its intended application; this is the approach taken, e.g., in [10, 25].

Algorithm 1 EM with latent disturbances

1. Set $k = 0$ and initialize θ_k such that $L_{\theta_k}(y_{1:T})$ is finite.
2. **Expectation (E) Step:**
   (2.1) Compute $\hat{x}_{1|T}$ and $\widehat{\Sigma}_{1|T}$ as in (14).
   (2.2) Compute $\widehat{\Sigma}_w$ as in (21).
   (2.3) Compute $\hat{Z}$ and Ω as in (15).
3. **Maximization (M) Step:**
   (3.1) Set $\alpha_{k+1} = \{\hat{x}_{1|T}, \widehat{\Sigma}_{1|T}\}$ and $\beta_{k+1} = \widehat{\Sigma}_w$.
   (3.2) Assemble $\{\lambda^i\}_{j=0}^T$ of the form (28) by computing $\{H_j\}_{j=0}^T$ with (34) and $\{h_j\}_{j=0}^T$ with (35).
   (3.3) Compute η_{k+1} by solving (36) and set $\gamma_{k+1} = \mathcal{M}(\eta_{k+1})$.
   (3.4) Set $\theta_{k+1} = \{\alpha_{k+1}, \beta_{k+1}, \gamma_{k+1}\}$.
4. Terminate if (37), otherwise $k \leftarrow k + 1$ and return to step 2.

5 Theoretical properties of identification via EM

5.1 Singular state space models

In applications, it may arise that the dimension of the disturbance is less than that of the state variable, i.e. $n_w < n_x$. For example, consider a simple mass-spring-damper system governed by $m\ddot{s} + c\dot{s} + ks = u + w$ for displacement s. When discretized, these dynamics can be represented by the second order state space model

$$x_{t+1} = \begin{bmatrix} 1 & \Delta_t \\ \frac{-k\Delta_t}{m} & 1 - \frac{c\Delta_t}{m} \end{bmatrix} x_t + \begin{bmatrix} 0 \\ \Delta_t \end{bmatrix} u + \begin{bmatrix} 0 \\ \Delta_t \end{bmatrix} w$$

with state variable $x_t = [s(t) \ \dot{s}(t)]'$.

In such cases, the process noise covariance $G\Sigma_v G'$ is singular, and standard EM algorithms based on latent states are no longer applicable. To see why, observe that the transition density of such a model is given by

$$p_\theta(x_{t+1} \mid x_t) = \mathcal{N}(x_{t+1}; Ax_t + Bu_t, G\Sigma_w G').$$

As $G\Sigma_w G'$ is rank deficient, the transition $p_\theta(x_{t+1} \mid x_t)$ does not admit a density [26], and so we cannot evaluate, much less optimize, the joint log likelihood $\log p_\theta(y_{1:T}, x_{1:T})$ given in (7). Modifications to the standard latent states EM algorithm have been proposed to circumvent this difficulty e.g. the work of [27] introduces a perturbation model with full-rank process noise covariance. However, by choosing latent disturbances we can elegantly handle identification of both singular and full-rank state space models, with the same algorithm.

In particular, when formulating the EM algorithm over latent disturbances we work with the joint likelihood function $p_\theta(y_{1:T}, x_1, w_{1:T})$, given in (8). Comparing (8) to (7), we observe that the problematic transition density is replaced by the joint distribution of disturbances

$$p_\theta(w_{1:T}) = \prod_{t=1}^{T} \mathcal{N}(w_t; 0, \Sigma_w).$$

This distribution is independent of n_x, and so $p_\theta(y_{1:T}, x_1, w_{1:T})$ and, therefore, $Q(\theta, \theta_k)$ remains well-defined, even in the singular case, $n_w < n_x$.

5.2 Absence of disturbances or output noise

In this section, we study the auxiliary function $Q(\theta, \theta_k)$ in the limit cases of $\Sigma_w = 0$ and $\Sigma_v = 0$, for different choices of latent variables. These results will offer insight into the behavior of the EM algorithm as a function of disturbance magnitude, as explored in the numerical experiments of Section 6.1. For convenience, we denote the bounds based on latent states and disturbances by $Q_{\text{ls}}(\theta, \theta_k)$ and $Q_{\text{ld}}(\theta, \theta_k)$, respectively.

Proposition 9. *Consider a model of the form (1), and let θ be such that $\Sigma_w = 0$, i.e. disturbances are omitted from the model. The auxiliary function built on* latent states, *$Q_{\text{ls}}(\theta, \theta_k)$, is undefined when $A \neq A_k$ or $B \neq B_k$.*

Proof. When $\Sigma_w = 0$, given any $x_1 \in \mathbb{R}^{n_x}$ the p.d.f. $p_{\theta_k}(x_{1:T} \mid y_{1:T})$ is nonzero on the set $\mathcal{S}(\theta_k) = \{x_{1:T} : x_{1:T} = \mathcal{X}(\theta_k, u_{1:T}, x_1, 0) \ \forall \ x_1 \in \mathbb{R}^{n_x}\}$. The auxiliary function may be expressed as

$$Q_{\text{ls}}(\theta, \theta_k) = \int_{\mathcal{S}(\theta_k)} \log p_\theta(x_{1:T}, y_{1:T}) p_{\theta_k}(x_{1:T} \mid y_{1:T}) \, \mathrm{d}x_{1:T}.$$

As $\Sigma_w = 0$, $p_\theta(x_{2:T} \mid x_1)$ is deterministic, evaluating to unity when $x_{1:T} = \mathcal{X}(\theta, u_{1:T}, x_1, 0)$, and zero otherwise. When $A \neq A_k$ or $B \neq B_k$, $\log p_\theta(x_{2:T} \mid x_1) = 0$ for all $x_{1:T} \in \mathcal{S}(\theta_k)$, and so $\log p_\theta(x_{1:T}, y_{1:T})$ is undefined. As a consequence, $Q_{\text{ls}}(\theta, \theta_k)$ is undefined.

When $A = A_k$ and $B = B_k$, $p_\theta(x_{2:T} \mid x_1) = 1$ for all $x \in \mathcal{S}(\theta_k)$ and so $Q_{\text{ls}}(\theta, \theta_k)$ can be evaluated as usual. ■

Proposition 10. *Consider a model of the form (1), and let θ be such that $\Sigma_w = 0$, i.e. disturbances are omitted from the model. Furthermore, suppose $\Sigma_1 = 0$; i.e. the initial conditions $x_1 = \mu$ are modeled without uncertainty. Then $L_\theta(y_{1:T}, x_1) = Q_{\text{ld}}(\theta, \theta_k)$ for all θ, θ_k; i.e., the auxiliary function built on* latent disturbances, *$Q_{\text{ld}}(\theta, \theta_k)$, reduces to the log likelihood.*

Proof. As $\Sigma_w = 0$, $\Sigma_1 = 0$ the p.d.f. $p_{\theta_k}(x_1, w_{1:T} \mid y_{1:T})$ is trivially deterministic, evaluating to unity when $x_1 = \mu$ and $w_{1:T} \equiv 0$, and evaluating to zero otherwise. Therefore

$$Q_{\text{ld}}(\theta, \theta_k) = \log p_\theta(y_{1:T}, \mu, 0) = \log p_\theta(y_{1:T} \mid \mu).$$

The log likelihood can be decomposed as

$$\begin{aligned} L_\theta(y_{1:T}) &= \log \int p_\theta(y_{1:T}, x_1) \mathrm{d}x_1 \\ &= \log \int p_\theta(y_{1:T} \mid x_1) p_\theta(x_1) \mathrm{d}x_1 = \log p_\theta(y_{1:T} \mid \mu), \end{aligned}$$

where the final equality follows from the fact that $p_\theta(x_1)$ is a δ-function, at $x_1 = \mu$. ■

Proposition 11. *Consider a* first order *model of the form (1), and let θ be such that $\Sigma_v = 0$, i.e. output noise is omitted from the model. The auxiliary function built on* latent disturbances, *$Q_{\text{ld}}(\theta, \theta_k)$, is undefined for $\theta \neq \theta_k$, i.e. $Q(\theta, \theta_k)$ collapses to a single point at $\theta = \theta_k$.*

Proof. For a given θ, let $x^\theta_{1:T}$ denote the unique state sequence that is 'consistent' with the data, i.e. $x^\theta_{1:T} \triangleq \{x_{1:T} : y_t = Cx_t + Du_t, t = 1, \ldots, T\}$. There is also a corresponding unique disturbance sequence, denoted $w^\theta_{1:T} = \{w_{1:T} : x^\theta_{1:T} = \mathcal{X}(\theta, u_{1:T}, x^\theta_1, w_{1:T})\}$.

As $\Sigma_v = 0$, the p.d.f. $p_{\theta_k}(x_1, w_{1:T} \mid y_{1:T})$ is a δ-function at $x_1 = x_1^{\theta_k}$ and $w_{1:T} = w_{1:T}^{\theta_k}$. The auxiliary function is then given by

$$Q_{\mathrm{ld}}(\theta, \theta_k) = \log p_\theta(y_{1:T}, x_1^{\theta_k}, w_{1:T}^{\theta_k}).$$

We can decompose $p_\theta(y_{1:T}, x_1^{\theta_k}, w_{1:T}^{\theta_k})$ as in (8). As $\Sigma_v = 0$, the p.d.f. $p_\theta(y_{1:T} \mid x_1, w_{1:T})$ is also a δ-function at $x_1 = x_1^\theta$ and $w_{1:T} = w_{1:T}^\theta$. If $C \neq C_k$ or $D \neq D_k$ then $x_1^\theta \neq x_1^{\theta_k}$. Furthermore, if $A \neq A_k$, $B \neq B_k$ or $G \neq G_k$, then $\mathcal{X}(\theta, u_{1:T}, x_1^\theta, w_{1:T}^\theta) \neq \mathcal{X}(\theta_k, u_{1:T}, x_1^{\theta_k}, w_{1:T}^{\theta_k})$. In both cases $p_\theta(y_{1:T} \mid x_1^{\theta_k}, w_{1:T}^{\theta_k}) = 0$ and so $Q_{\mathrm{ld}}(\theta, \theta_k)$ is undefined.

When $\theta = \theta_k$, $p_\theta(y_{1:T} \mid x_1^{\theta_k}, w_{1:T}^{\theta_k}) = 1$ and $Q_{\mathrm{ld}}(\theta, \theta_k)$ can be evaluated as usual. ■

Proposition 12. *Consider a* first order *model of the form (1), and let θ be such that $\Sigma_v = 0$, i.e. output noise is omitted from the model. Let $Q_{\mathrm{ls}}(\theta, \theta_k)$ denote the auxiliary function built on* latent states*, then:*

i. $Q_{\mathrm{ls}}(\theta, \theta_k)$ is undefined for all θ such that $C \neq C_k$ or $D \neq D_k$.

ii. $Q_{\mathrm{ls}}(\theta, \theta_k) = L_\theta(y_{1:T})$ for all θ such that $C = C_k$ and $D = D_k$.

Proof. For a given θ, let $x_{1:T}^\theta$ denote the unique state sequence that is 'consistent' with the data, i.e. $x_{1:T}^\theta \triangleq \{x_{1:T} \colon y_t = Cx_t + Du_t, t = 1, \ldots, T\}$. As $\Sigma_v = 0$, given $y_{1:T}$ both $p_{\theta_k}(x_{1:T} \mid y_{1:T})$ and $p_\theta(y_{1:T} \mid x_{1:T})$ are δ-functions at $x_{1:T} = x_{1:T}^\theta$. The auxiliary function is then given by

$$Q_{\mathrm{ls}}(\theta, \theta_k) = \log p_\theta(y_{1:T}, x_{1:T}^{\theta_k}).$$

Let us now consider the two cases:

i. When $C \neq C_k$ or $D \neq D_k$, $x_{1:T}^\theta \neq x_{1:T}^{\theta_k}$ and so $p_\theta(y_{1:T} \mid x_{1:T}^{\theta_k}) = 0$. Therefore, $Q_{\mathrm{ls}}(\theta, \theta_k)$ is undefined.

ii. When $C = C_k$ and $D = D_k$, $x_{1:T}^\theta = x_{1:T}^{\theta_k}$ and so

$$Q_{\mathrm{ls}}(\theta, \theta_k) = \log p_\theta(y_{1:T} \mid x_{1:T}^\theta) p_\theta(x_{1:T}^\theta) = \log p_\theta(x_{1:T}^\theta).$$

The likelihood can be expressed as

$$L_\theta(y_{1:T}) = \log \int p_\theta(y_{1:T} \mid x_{1:T}) p_\theta(x_{1:T}) \mathrm{d}x_{1:T} = \log p_\theta(x_{1:T}^\theta),$$

where the second inequality comes from the fact that $p_\theta(y_{1:T} \mid x_{1:T})$ is a δ-function. Therefore, $L_\theta(y_{1:T}) = Q_{\mathrm{ls}}(\theta, \theta_k)$.

■

6 Numerical experiments

6.1 Influence of disturbance magnitude on bound fidelity

In the following experiment, we investigate the fidelity of $Q(\theta, \theta_k)$ as a bound on $L_\theta(y_{1:T})$, as a function of the magnitude of the disturbances, $w_{1:T}$, and the choice of latent variables. As in Section 5.2, we denote the bounds based on latent states and disturbances by $Q_{\mathrm{ls}}(\theta, \theta_k)$ and $Q_{\mathrm{ld}}(\theta, \theta_k)$, respectively. The results are presented in Figure 1, which depicts Q_{ls}, Q_{ld} and $L_\theta(y_{1:T})$ for a first order ($n_x = 1$) LGSS model, each plotted as a function of the single unknown scalar parameter $\theta = A$.

We begin with the case of 'small' disturbances (i.e. $\Sigma_w \ll \Sigma_v$) as depicted in Figure 1(a), and observe the following: $Q_{\mathrm{ld}}(\theta, \theta_k)$ represents $L_\theta(y_{1:T})$ with high fidelity, whereas $Q_{\mathrm{ls}}(\theta, \theta_k)$ is localized about θ_k. Such an observation is not without precedent. For instance, in the latent states formulation of [9, Section 10] it was noted that an initial disturbance covariance estimate $\Sigma_w = 0$ results in $\theta_k = \theta_0$ for all k; i.e. the model parameters are not improved. This suggests that $Q_{\mathrm{ls}}(\theta, \theta_k)$ fails to accurately represent $L_\theta(y_{1:T})$, except at $\theta = \theta_0$.

Proposition 9 makes this observation more precise: in the 1D case of Figure 1(a), when $\Sigma_w = 0$, $Q_{\mathrm{ls}}(\theta, \theta_k)$ is undefined for $A \neq A_k$. Taken together, Figure 1(a) and Proposition 9 suggest that as Σ_w becomes smaller (relative to Σ_v) the bound $Q_{\mathrm{ls}}(\theta, \theta_k)$ becomes more localized about θ_k, eventually collapsing to a single point when $\Sigma_w = 0$. Conversely, as Σ_w (and Σ_1) decrease, $Q_{\mathrm{ld}}(\theta, \theta_k)$ becomes an increasingly accurate representation of the log likelihood, eventually reproducing $L_\theta(y_{1:T})$ *exactly*, when Σ_w (and Σ_1) are identically zero, as in Proposition 10.

Turning our attention to the case of 'large' disturbances (i.e. $\Sigma_w \gg \Sigma_v$) as depicted in Figure 1(b), we observe the opposite behavior: $Q_{\mathrm{ls}}(\theta, \theta_k)$ faithfully represents the log likelihood, whereas $Q_{\mathrm{ld}}(\theta, \theta_k)$ appears to be localized about θ_k. Once more, studying the limiting case $\Sigma_v = 0$ offers insight into this behavior: Proposition 11 states that when $\Sigma_v = 0$, $Q_{\mathrm{ld}}(\theta, \theta_k)$ is undefined for $A \neq A_k$.

Taken together, Figure 1(b) and Proposition 11 suggest that as Σ_v decreases (i.e. as Σ_w increases relative to Σ_v), the bound $Q_{\mathrm{ld}}(\theta, \theta_k)$ becomes more localized about θ_k, eventually collapsing to a single point when $\Sigma_v = 0$. Conversely, for this 1D experiment with $\theta = A$, Proposition 12 states that $Q_{\mathrm{ls}}(\theta, \theta_k)$ will reproduce $L_\theta(y_{1:T})$ *exactly*, when Σ_v is identically zero. Indeed, in Figure 1(b) with $\Sigma_v \ll \Sigma_w$, we observe $Q_{\mathrm{ls}}(\theta, \theta_k)$ representing the likelihood faithfully.

To summarize: in the case of 'large disturbances' (i.e. $\Sigma_w \gg \Sigma_v$), $Q_{\mathrm{ld}}(\theta, \theta_k)$ will tend to bound $L_\theta(y_{1:T})$ with greater fidelity, compared to $Q_{\mathrm{ls}}(\theta, \theta_k)$. In the case of 'small disturbances' (i.e. $\Sigma_w \ll \Sigma_v$) the converse is true.

6.2 Convergence rate

It is clear from Figure 1(a), that both $Q_{\mathrm{ld}}(\theta, \theta_k)$ and $\bar{Q}_3(\eta)$ better represent $L_\theta(y_{1:T})$ compared to $Q_{\mathrm{ls}}(\theta, \theta_k)$. In fact, one would expect that optimization of $\bar{Q}_3(\eta)$, as in Algorithm 1, would converge to θ^{ML} in fewer iterations than optimization of $Q_{\mathrm{ls}}(\theta, \theta_k)$, as in a 'standard' EM algorithm [10].

This principle, which is clearly understood in the first order example of Figure 1, is further illustrated in Figure 3 for three different 4^{th} order SISO systems; Bode plots for each system are given in Figure 2. The results in Figure 3 clearly show Algorithm 1, based on latent disturbances, converging in fewer iterations than the latent states formulation of [10]. These results are consistent with the analysis in Section 6.1. Specifically, in each trial disturbances were 'small' in magnitude ($\Sigma_w = 1 \times 10^{-5}$) and so we expect $\bar{Q}_3(\eta)$ to better represent the likelihood, allowing Algorithm 1 to converge in fewer iterations.

It should be stressed that although Algorithm 1 converges in fewer iterations than the latent states formulation, each iteration is considerably more computationally expensive, and thus the total computation times for each algorithm are comparable. Nevertheless, this faster convergence rate is advantageous as it renders Algorithm 1 less sensitive to the choice of δ when termination conditions of the form (37) are employed. Furthermore, as methods for SDP mature, one may expect Algorithm 1 to gain the upper hand in regards to computation time.

6.3 Stability of the identified model

A desirable property of Algorithm 1 is that stability of the identified model is enforced at every iteration; recall from Lemma 6 that we confine our search to an implicit parametrization of all stable models, which, by Lemma 7, is necessary to ensure that $\bar{J}_\lambda(\eta)$ is well-defined. Conversely, in a standard latent states implementation of the EM algorithm, the M step is accomplished by the solution of an unconstrained linear least squares problem [10]. Consequently, it is possible that at any iteration (or indeed the conclusion) of the algorithm, the parameters θ_k could constitute an unstable model.

Such a scenario is illustrated in the numerical experiment of Figure 4, which depicts the identification of a 4^{th} order model, similar to System 2 in Figure 2. From Figure 4(b) it is apparent that, for the first one thousand iterations, the parameters maintained by the latent states EM algorithm represent an *unstable* model (i.e. $|\lambda_{\max}(A_k)| > 1$). This instability can be particularly

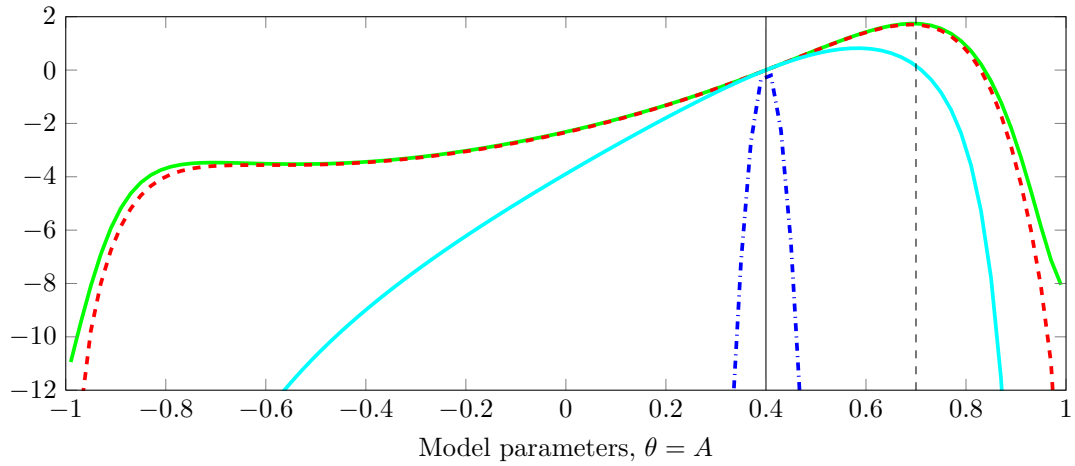

(a) 'Small' disturbances: $\Sigma_v = 1 \times 10^{-3}$ and $\Sigma_v = 1 \times 10^{-2}$.

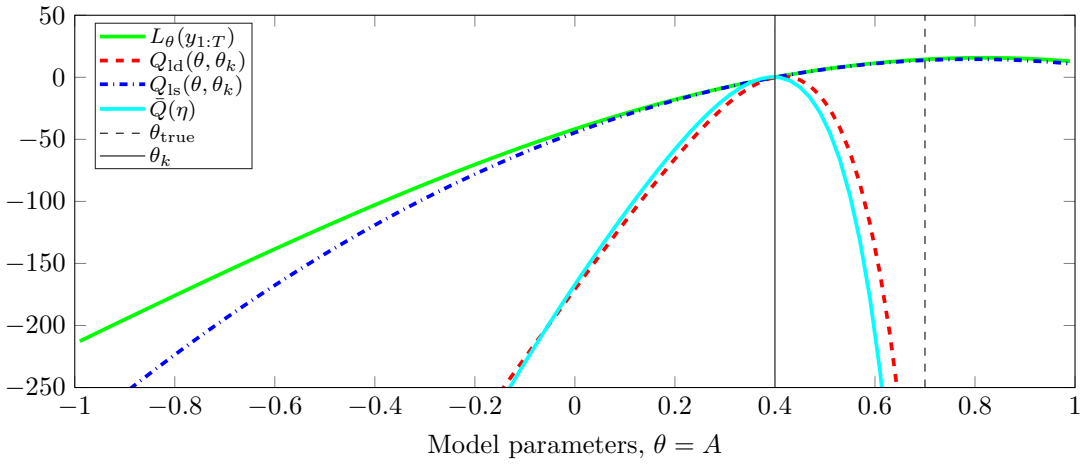

(b) 'Large' disturbances: $\Sigma_w = 10$ and $\Sigma_v = 1 \times 10^{-2}$.

Figure 1: Lower bounds to the log likelihood $L_\theta(y_{1:T})$ of a first order system with a single unknown scalar parameter, A. $Q_{\mathrm{ld}}(\theta, \theta_k)$ and $Q_{\mathrm{ls}}(\theta, \theta_k)$ denote the bounds based on latent disturbances and states respectively, while $\bar{Q}(\eta) = Q_1(\alpha, \theta_k) + Q_2(\beta, \theta_k) + \bar{Q}_3(\eta)$, where $\bar{Q}_3(\eta)$ is the bound based on Lagrangian relaxation defined in (27).

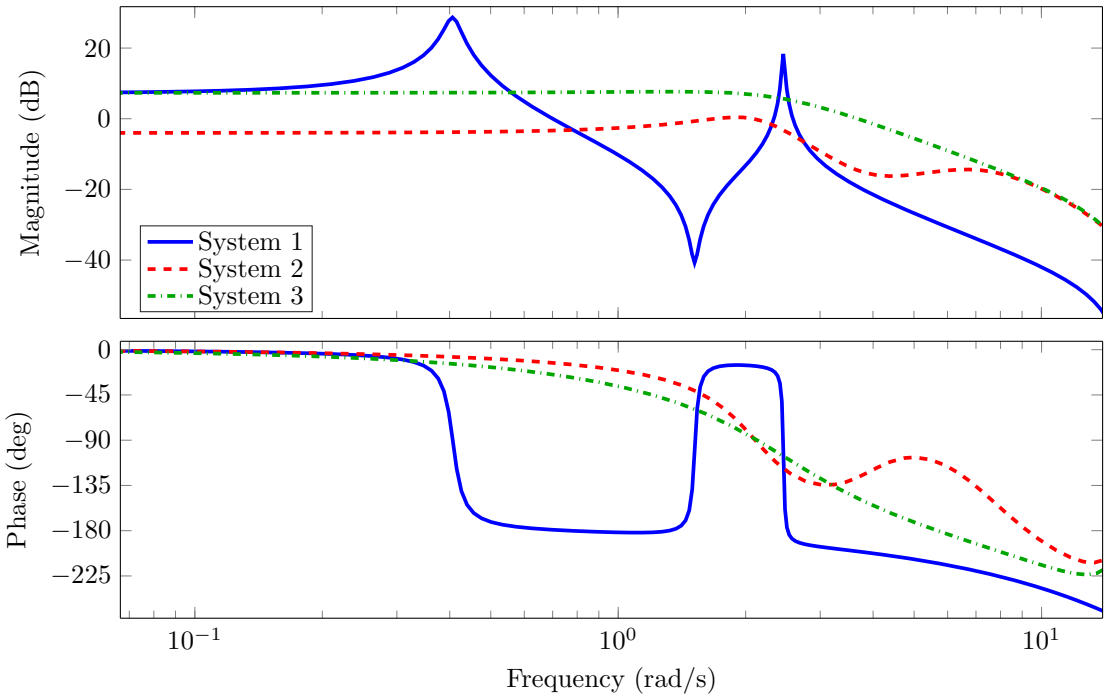

Figure 2: Bode plots of 4^{th} order systems used for the experiments presented in Figure 3.

problematic, given the slow convergence rate; e.g. in this instance, if a heuristic such as (37) was used employed, for $\delta > 4.7 \times 10^{-3}$ the algorithm would terminate before the thousandth iteration, and an unstable model would be returned. Conversely, the parameters maintained by Algorithm 1 constitute a stable model at each iteration.

7 Conclusion

In this paper, we have formulated the EM algorithm over latent disturbances, rather than states, for the identification of linear dynamical systems. Our main contribution is the use of Lagrangian relaxation to obtain a convex approximation of the challenging maximization step, guaranteed not to decrease the likelihood at each iteration. Though more computationally complex, this formulation with latent disturbances allows EM to be applied to singular state-space models, where latent states based methods break down.

Extension of this approach to the identification of nonlinear models shall be the subject of future research. In the nonlinear case, two major challenges arise during the formulation of EM with latent disturbances. First, the E step (c.f. Section 3.1) now involves a nonlinear disturbance smoothing problem, for which no closed form solution is known to exist. In recent decades, *sequential Monte Carlo* (SMC) methods [28] have emerged as effective tools for overcoming similar difficulties, having already proved useful in nonlinear, non-Gaussian *state smoothing* [29] and *disturbance filtering* [30] problems.

Second, nonlinearity of the model complicates the Lagrangian relaxation of the M step; e.g. the bound $\hat{J}_\lambda(\eta)$ cannot be evaluated analytically, as the supremum (in (24)) requires optimization of a function that is no longer quadratic in x. To proceed, one might approximate the simulation error terms in $Q_3(\gamma, \theta_k)$ with the *linearized simulation error*, introduced in [12], to which the Lagrangian relaxation presented in this work can be applied with little modification. Alternatively, when the system nonlinearity is modeled as a polynomial, *sum-of-squares* (SOS) programing [31] may be used to generate, and optimize, convex approximations to the Lagrangian relaxation.

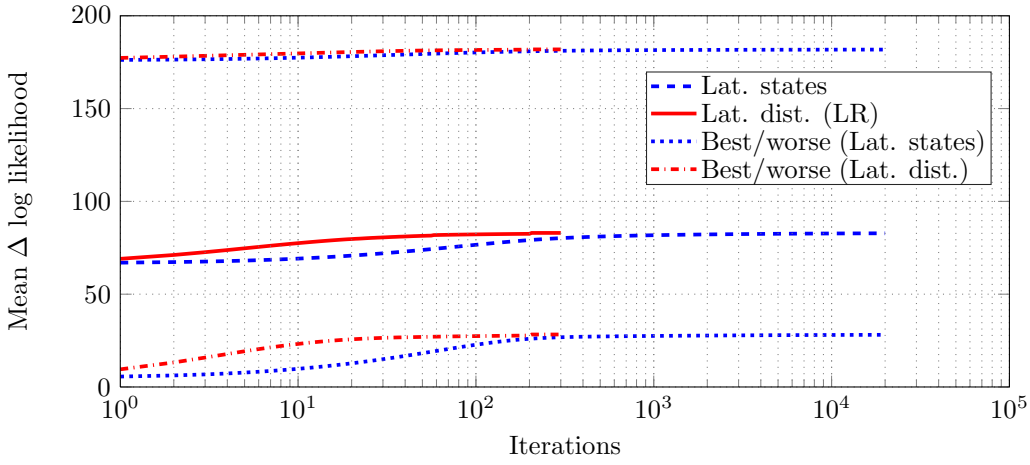

(a) System 1, sharp resonant peaks.

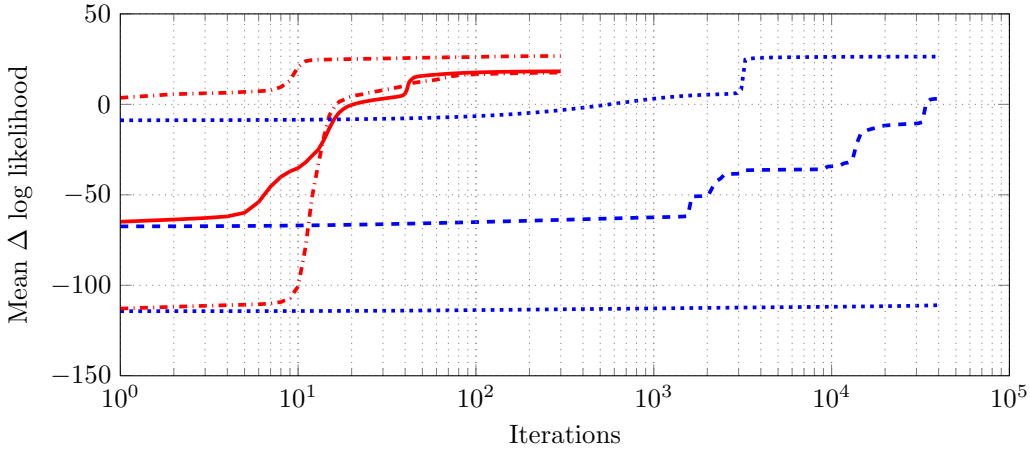

(b) System 2, smooth resonant peaks.

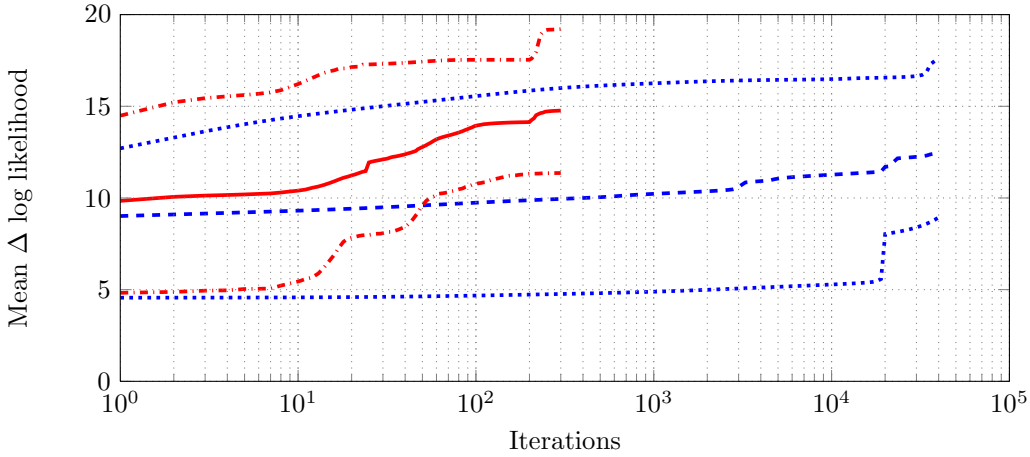

(c) System 3, overdamped.

Figure 3: Difference between $L_{\theta_k}(y_{1:T})$ and $L_{\theta_{\text{true}}}(y_{1:T})$ (where θ_{true} denotes the true model parameters) as a function of iterations for Algorithm 1 (EM with latent disturbances, red) and the method of [10] (EM with latent states, blue dashed). The difference is averaged over 10 trials, each with SNR of 100, $\Sigma_v = 1 \times 10^{-5}$ and $T = 250$. Also plotted are the best and worst trial results (in terms of final likelihood) for each system. Bode plots for systems 1, 2 and 4 are depicted in Figure 2.

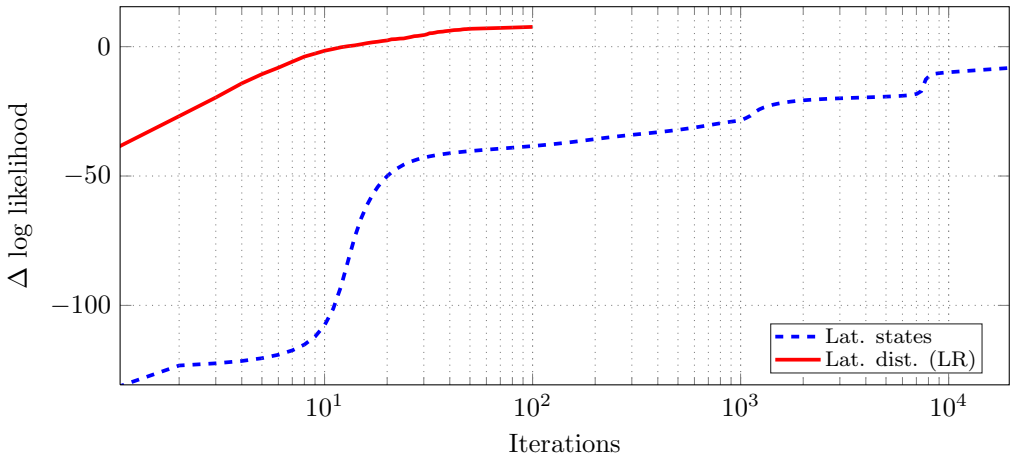

(a) Difference between $L_{\theta_k}(y_{1:T})$ and $L_{\theta_{\text{true}}}(y_{1:T})$, where θ_{true} denotes the true model parameters, at each iteration.

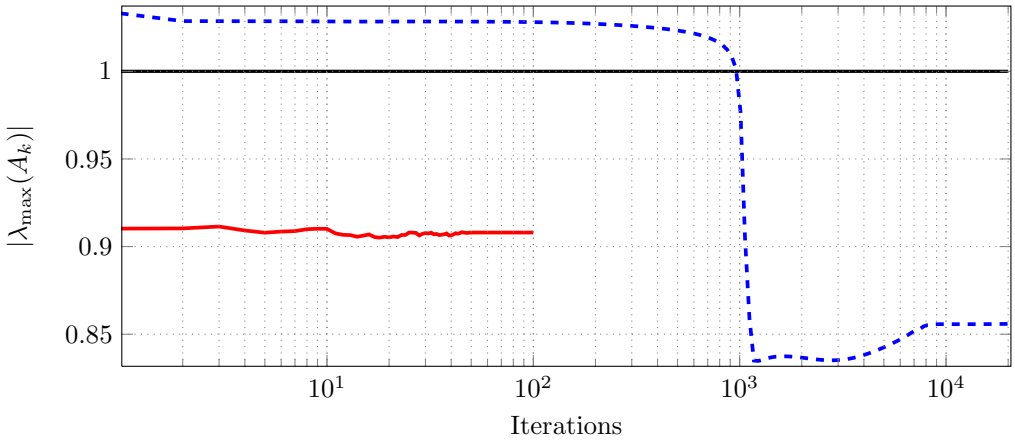

(b) Magnitude of the largest eigenvalue of A_k, at each iteration. When the spectral radius of A_k is greater than unity, i.e. $|\lambda_{\max}(A_k)| > 1$, the model θ_k is unstable.

Figure 4: Log likelihood and spectral radius of A_k at each iteration for two different EM algorithms: i. EM with latent states as in [10] (Lat. states, blue dash); ii. Algorithm 1 (Lat. dist. (LR), red). The spectral radius of A for the true system was 0.90.

References

[1] L. Ljung, *System Identification: Theory for the User (2nd Edition)*, 2nd ed. Prentice Hall, Jan. 1999.

[2] ——, "Prediction error estimation methods," *Circuits, Systems and Signal Processing*, vol. 21, no. 1, pp. 11–21, 2002.

[3] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, "Equation of state calculations by fast computing machines," *The journal of chemical physics*, vol. 21, no. 6, pp. 1087–1092, 1953.

[4] W. K. Hastings, "Monte Carlo sampling methods using Markov chains and their applications," *Biometrika*, vol. 57, no. 1, pp. 97–109, 1970.

[5] L. Vandenberghe and S. Boyd, "Semidefinite programming," *SIAM review*, vol. 38, no. 1, pp. 49–95, 1996. [Online]. Available: http://epubs.siam.org/doi/abs/10.1137/1038003

[6] W. E. Larimore, "System identification, reduced-order filtering and modeling via canonical variate analysis," in *American Control Conference, 1983*. IEEE, 1983, pp. 445–451.

[7] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 6, no. 6, pp. 721–741, Nov. 1984.

[8] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *Journal of the royal statistical society. Series B (methodological)*, pp. 1–38, 1977.

[9] T. B. Schön, A. Wills, and B. Ninness, "System identification of nonlinear state-space models," *Automatica*, vol. 47, no. 1, pp. 39–49, 2011.

[10] S. Gibson and B. Ninness, "Robust maximum-likelihood estimation of multivariable dynamic systems," *Automatica*, vol. 41, no. 10, pp. 1667–1682, 2005.

[11] R. H. Shumway and D. S. Stoffer, "An approach to time series smoothing and forecasting using the EM algorithm," *Journal of time series analysis*, vol. 3, no. 4, pp. 253–264, 1982.

[12] M. M. Tobenkin, I. R. Manchester, J. Wang, A. Megretski, and R. Tedrake, "Convex optimization in identification of stable non-linear state space models," in *Proceedings of the 49th IEEE Conference on Decision and Control, CDC 2010, December 15-17, 2010, Atlanta, Georgia, USA*. IEEE, 2010, pp. 7232–7237.

[13] B. N. B. Bond, Z. Mahmood, Y. Li, R. Sredojevic, A. Megretski, V. Stojanovi, Y. Avniel, and L. Daniel, "Compact modeling of nonlinear analog circuits using system identification via semidefinite programming and incremental stability certification," *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, vol. 29, no. 8, pp. 1149–1162, 2010.

[14] A. Megretski, "Convex optimization in robust identification of nonlinear feedback," in *Proceedings of the 47th IEEE Conference on Decision and Control, CDC*. IEEE, 2008, pp. 1370–1374.

[15] C. Lemaréchal, "Lagrangian relaxation," in *Computational combinatorial optimization*. Springer, 2001, pp. 112–156.

[16] V. A. Yakubovich, "S-procedure in nonlinear control theory," *Vestnik Leningrad University*, vol. 1, pp. 62–77, 1971.

[17] I. Pólik and T. Terlaky, "A survey of the s-lemma," *SIAM review*, vol. 49, no. 3, pp. 371–418, 2007.

[18] J. M. Ortega and W. C. Rheinboldt, *Iterative solution of nonlinear equations in several variables*. Siam, 1970, vol. 30.

[19] D. R. Hunter and K. Lange, "A tutorial on MM algorithms," *The American Statistician*, vol. 58, no. 1, pp. 30–37, 2004.

[20] J. Umenberger, J. Wågberg, I. Manchester, and T. B. Schön, "On identification via EM with latent disturbances and Lagrangian relaxation," in *Proceedings of the 17th IFAC Symposium on System Identification (SYSID)*, Beijing, China, 2015.

[21] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge: Cambridge University Press, 2004.

[22] H. E. Rauch, C. Striebel, and F. Tung, "Maximum likelihood estimates of linear dynamic systems," *AIAA journal*, vol. 3, no. 8, pp. 1445–1450, 1965.

[23] J. Durbin and S. J. Koopman, *Time series analysis by state space methods*, ser. Oxford Statistical Science Series. Oxford University Press, 2012, no. 38.

[24] I. Manchester, M. M. Tobenkin, and A. Megretski, "Stable nonlinear system identification: Convexity, model class, and consistency," in *Proceedings of the 16th IFAC Symposium on System Identification (SYSID)*, Brussels, Belgium, 2012.

[25] A. Wills, T. B. Schön, and B. Ninness, "Estimating state-space models in innovations form using the expectation maximisation algorithm," in *Proceedings of the 49th IEEE Conference on Decision and Control (CDC)*. IEEE, 2010, pp. 5524–5529.

[26] K. E. Muller and P. W. Stewart, *Linear Model Theory: Univariate, Multivariate, and Mixed Models*. John Wiley & Sons, Inc., 2006.

[27] V. Solo, "An EM algorithm for singular state space models," in *Proceedings of the 42nd IEEE Conference on Decision and Control, CDC 2003.*, vol. 4. IEEE, 2003, pp. 3457–3460.

[28] N. J. Gordon, D. J. Salmond, and A. F. Smith, "Novel approach to nonlinear/non-Gaussian Bayesian state estimation," *IEEE Proceedings F, Radar and Signal Processing*, vol. 140, no. 2, pp. 107–113, Apr. 1993.

[29] T. B. Schön, F. Lindsten, J. Dahlin, J. Wågberg, C. A. Naesseth, A. Svensson, and L. Dai, "Sequential Monte Carlo methods for system identification," *IFAC-PapersOnLine*, vol. 48, no. 28, pp. 775–786, 2015.

[30] L. M. Murray, E. M. Jones, and J. Parslow, "On disturbance state-space models and the particle marginal Metropolis-Hastings sampler," *SIAM/ASA Journal on Uncertainty Quantification*, vol. 1, no. 1, pp. 494–521, 2013.

[31] P. A. Parrilo, "Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization," Ph.D. dissertation, California Institute of Technology, May 2000.